

Discrete mean field games: the short-stage limit

Juan Pablo Maldonado López

Equipe Combinatoire et Optimisation, Université Pierre et Marie Curie, Paris

Tours, June 26, 2014

Outline

- 1 Introduction
- 2 The N-player game
- 3 Mean field equilibrium
- 4 The short-stage model

Mean field games

Mean field games models aim to understand the behavior of a large number of identical players, where each tries to optimize its position in space and time, but whose preferences are determined by the choices of the other players.

Mean field games

Mean field games models aim to understand the behavior of a large number of identical players, where each tries to optimize its position in space and time, but whose preferences are determined by the choices of the other players.

- Introduced by Huang, Caines and Malhamé (2003, 2006) and by Lasry and Lions (2006,2007).

Mean field games

Mean field games models aim to understand the behavior of a large number of identical players, where each tries to optimize its position in space and time, but whose preferences are determined by the choices of the other players.

- Introduced by Huang, Caines and Malhamé (2003, 2006) and by Lasry and Lions (2006,2007).
- Two important features: dynamics and anticipation (backward-forward structure.)

Mean field games

Mean field games models aim to understand the behavior of a large number of identical players, where each tries to optimize its position in space and time, but whose preferences are determined by the choices of the other players.

- Introduced by Huang, Caines and Malhamé (2003, 2006) and by Lasry and Lions (2006,2007).
- Two important features: dynamics and anticipation (backward-forward structure.)
- Most of the studied models are in continuous time: the "backward" part corresponding to a Hamilton-Jacobi PDE and the "forward" part corresponding to a Fokker-Planck PDE.

Why discrete time?

- Easier to handle: no PDE's.

Why discrete time?

- Easier to handle: no PDE's.
- Possible to choose actions randomly: in continuous time, players have a unique optimal action at each stage.

Why discrete time?

- Easier to handle: no PDE's.
- Possible to choose actions randomly: in continuous time, players have a unique optimal action at each stage.
- Gomes, Mohr and Souza (2010) introduced a model for discrete time mean field games, but they consider a continuum of players instead and study asymptotic behaviour of the game as the time horizon tends to infinity. Our asymptotic results concern the number of players.

Why discrete time?

- Easier to handle: no PDE's.
- Possible to choose actions randomly: in continuous time, players have a unique optimal action at each stage.
- Gomes, Mohr and Souza (2010) introduced a model for discrete time mean field games, but they consider a continuum of players instead and study asymptotic behaviour of the game as the time horizon tends to infinity. Our asymptotic results concern the number of players.
- Our model builds on results from Adlakha, Johari and Weintraub (2012).

Notation

- Let \mathcal{X} denote the state space and A the action set (both finite).

Notation

- Let \mathcal{X} denote the state space and A the action set (both finite).
- The payoff functions:

$$\ell : \mathcal{X} \times A \times \mathcal{P}(\mathcal{X}) \rightarrow [0, 1], \quad g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$$

Notation

- Let \mathcal{X} denote the state space and A the action set (both finite).
- The payoff functions:

$$\ell : \mathcal{X} \times A \times \mathcal{P}(\mathcal{X}) \rightarrow [0, 1], \quad g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$$

- The transition function:

$$Q : \mathcal{X} \times A \times \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$$

The game with finitely many players

- Let I denote the set of players and assume $|I| = N$.

The game with finitely many players

- Let I denote the set of players and assume $|I| = N$.
- Let $X_{t,N}^j$ be a random variable that describes the position of player j at time t .

The game with finitely many players

- Let I denote the set of players and assume $|I| = N$.
- Let $X_{t,N}^j$ be a random variable that describes the position of player j at time t .
- The states of the players at time $t = 0$ are chosen i.i.d using the distribution m_0 .

The game with finitely many players

- Let I denote the set of players and assume $|I| = N$.
- Let $X_{t,N}^j$ be a random variable that describes the position of player j at time t .
- The states of the players at time $t = 0$ are chosen i.i.d using the distribution m_0 .
- We reserve capital letters for random variables and lower case letters for their realizations.

Development of the game

- At stage $t = 0, 1 \dots T - 1$, player i observes his own state $x_{t,N}^i$ and the average position on the state space of all the players, $m_{t,N}$.

Development of the game

- At stage $t = 0, 1 \dots T - 1$, player i observes his own state $x_{t,N}^i$ and the average position on the state space of all the players, $m_{t,N}$.
- Players choose simultaneously and independently their actions $a_{t,N}^i$.

Development of the game

- At stage $t = 0, 1 \dots T - 1$, player i observes his own state $x_{t,N}^i$ and the average position on the state space of all the players, $m_{t,N}$.
- Players choose simultaneously and independently their actions $a_{t,N}^i$.
- Each player i receives the payoff $\ell(x_{t,N}^i, a_{t,N}^i, m_{t,N})$.

Development of the game

- At stage $t = 0, 1 \dots T - 1$, player i observes his own state $x_{t,N}^i$ and the average position on the state space of all the players, $m_{t,N}$.
- Players choose simultaneously and independently their actions $a_{t,N}^i$.
- Each player i receives the payoff $\ell(x_{t,N}^i, a_{t,N}^i, m_{t,N})$.
- The new state $X_{t+1,N}^i$ is chosen randomly using the transition function $Q(x_{t,N}^i, a_{t,N}^i, m_{t,N})$.

- The random average distribution is $M_{t+1,N} := \frac{1}{N} \sum_{j \neq i} \delta_{X_{t+1,N}^j}$.

- The random average distribution is $M_{t+1,N} := \frac{1}{N} \sum_{j \neq i} \delta_{X_{t+1,N}^j}$.
- At the beginning of stage $t + 1$, the realization of $X_{t+1,N}^i$ and $M_{t+1,N}$, denoted $x_{t+1,N}^i$ and $m_{t+1,N}$ respectively, are observed, and the situation is repeated.

- The random average distribution is $M_{t+1,N} := \frac{1}{N} \sum_{j \neq i} \delta_{X_{t+1,N}^j}$.
- At the beginning of stage $t + 1$, the realization of $X_{t+1,N}^i$ and $M_{t+1,N}$, denoted $x_{t+1,N}^i$ and $m_{t+1,N}$ respectively, are observed, and the situation is repeated.
- At stage $t = T$ a final payoff $g(x_{T,N}^i, m_{T,N})$ is allocated.

Strategies

- A **behavioral strategy** for player i is a vector $\pi^i = (\pi_t^i)_{t=0}^{T-1}$ where $\pi_t^i : \mathcal{H}_t \rightarrow \mathcal{P}(A)$ and $\mathcal{H}_t = (\mathcal{X} \times A \times \mathcal{P}(\mathcal{X}))^t$ is the set of all possible histories up to date t . Denote by \mathcal{S} the set of behavioral strategies.

Strategies

- A **behavioral strategy** for player i is a vector $\pi^i = (\pi_t^i)_{t=0}^{T-1}$ where $\pi_t^i : \mathcal{H}_t \rightarrow \mathcal{P}(A)$ and $\mathcal{H}_t = (\mathcal{X} \times A \times \mathcal{P}(\mathcal{X}))^t$ is the set of all possible histories up to date t . Denote by \mathcal{S} the set of behavioral strategies.
- A **Markovian strategy** for player i is a vector $\sigma^i = (\sigma_t^i)_{t=0}^{T-1}$ such that $\sigma_t^i : \mathcal{X} \rightarrow \mathcal{P}(A)$.
- A **stationary strategy** is a function $\sigma : \mathcal{X} \rightarrow \mathcal{P}(A)$.

Strategies

- A **behavioral strategy** for player i is a vector $\pi^i = (\pi_t^i)_{t=0}^{T-1}$ where $\pi_t^i : \mathcal{H}_t \rightarrow \mathcal{P}(A)$ and $\mathcal{H}_t = (\mathcal{X} \times A \times \mathcal{P}(\mathcal{X}))^t$ is the set of all possible histories up to date t . Denote by \mathcal{S} the set of behavioral strategies.
- A **Markovian strategy** for player i is a vector $\sigma^i = (\sigma_t^i)_{t=0}^{T-1}$ such that $\sigma_t^i : \mathcal{X} \rightarrow \mathcal{P}(A)$.
- A **stationary strategy** is a function $\sigma : \mathcal{X} \rightarrow \mathcal{P}(A)$.
- A **strategy profile** is a vector $\pi = (\pi^i)_{i \in I}$, where π^i is a behavioral strategy of player i .

Payoff

The payoff of player i , when using the strategy π^i and when his adversaries use the strategy profile $\pi^{-i} \in \mathcal{S}^{N-1}$ is

$$J_N^i(x, m_0, \pi^i, \pi^{-i}) := \mathbb{E}_\pi^Q \left[\sum_{t \in \mathcal{T}} \ell(x_{t,N}^i, a_{t,N}^i, m_{t,N}) + g(x_{T,N}^i, m_{T,N}) \right].$$

Nash equilibrium

Definition

An ϵ -**Nash equilibrium** where $\epsilon > 0$, is a strategy profile $(\pi^i)_{i \in I}$ such that, for all player i and all behavioral strategy τ ,

$$J_N^i(x, m_0, \tau, \pi^{-i}) \leq J_N^i(x, m_0, \pi^i, \pi^{-i}) + \epsilon.$$

How to play well in the N player game?

- The trick: Hide the complexity of the problem in a parameter $\mathbf{m} \in \mathcal{P}(\mathcal{X})$.

How to play well in the N player game?

- The trick: Hide the complexity of the problem in a parameter $\mathbf{m} \in \mathcal{P}(\mathcal{X})$.
- Anticipate a fixed average distribution on the state space $\mathbf{m} = (m_0, m_1, \dots, m_T)$.

How to play well in the N player game?

- The trick: Hide the complexity of the problem in a parameter $\mathbf{m} \in \mathcal{P}(\mathcal{X})$.
- Anticipate a fixed average distribution on the state space $\mathbf{m} = (m_0, m_1, \dots, m_T)$.
- Compute an optimal strategy for a single player with \mathbf{m} fixed.

How to play well in the N player game?

- The trick: Hide the complexity of the problem in a parameter $\mathbf{m} \in \mathcal{P}(\mathcal{X})$.
- Anticipate a fixed average distribution on the state space $\mathbf{m} = (m_0, m_1, \dots, m_T)$.
- Compute an optimal strategy for a single player with \mathbf{m} fixed.
- Everyone's strategy creates a vector of state distributions $\mathbf{m}' = (m_0, m'_1, \dots, m'_T)$

How to play well in the N player game?

- The trick: Hide the complexity of the problem in a parameter $\mathbf{m} \in \mathcal{P}(\mathcal{X})$.
- Anticipate a fixed average distribution on the state space $\mathbf{m} = (m_0, m_1, \dots, m_T)$.
- Compute an optimal strategy for a single player with \mathbf{m} fixed.
- Everyone's strategy creates a vector of state distributions $\mathbf{m}' = (m_0, m'_1, \dots, m'_T)$
- If $\mathbf{m} = \mathbf{m}'$, we are happy.

Dynamic programming

From the familiar arguments we obtain the following dynamic programming equation:

Dynamic programming

From the familiar arguments we obtain the following dynamic programming equation:

$$V(s, x, \mathbf{m}) = \max_{a \in A} \left\{ \ell(x, a, m_s) + \mathbb{E}^Q V(s+1, x_{s+1}, \mathbf{m}) \right\}$$

with terminal condition $V(T, x, \mathbf{m}) = g(x, m_T)$.

The forward component

Now consider a Markovian strategy $\sigma \in \Sigma$ and m_0 fixed and let $m_0^\sigma := m_0$. We define, for $t \geq 0$:

The forward component

Now consider a Markovian strategy $\sigma \in \Sigma$ and m_0 fixed and let $m_0^\sigma := m_0$. We define, for $t \geq 0$:

$$m_{t+1}^\sigma(x) := \sum_{y \in \mathcal{X}} Q(y, \sigma_t(y), m_t^\sigma)(x) \cdot m_t^\sigma(y).$$

Definition: MFE

Definition

Let m_0 given. A **mean field equilibrium** is a pair

$(\sigma, \mathbf{m}) = \left((\sigma_t)_{t=0}^{T-1}, (m_t)_{t=1}^T \right)$ such that:

- ① σ is the optimal strategy in the one player game $\Gamma_{\mathbf{m}}$, computed using the dynamic programming equation.
- ② \mathbf{m} is the trajectory followed by m_0 according to the mass equation for the strategy σ .

Existence

Proposition

(M.,2013) *There exists a mean field equilibrium for the finite horizon game in the following cases:*

- *If there exists a unique maximizer for the right hand side of dynamic programming equation for \mathbf{m} and for each (s, x) .*

Existence

Proposition

(M.,2013) *There exists a mean field equilibrium for the finite horizon game in the following cases:*

- *If there exists a unique maximizer for the right hand side of dynamic programming equation for \mathbf{m} and for each (s, x) .*
- *The transitions are independent of the state distribution, i.e.*

$$Q(x, a, m)(y) =: Q(x, a)(y), \quad \forall (x, y, a, m).$$

Approximation

Proposition

(M., 2013) Let x be a fixed initial state, (σ, \mathbf{m}) a mean field equilibrium and $(a_{t,N}^i)_{t \in \mathcal{T}}$ an arbitrary sequence of actions of player i . Consider the following two trajectories:

- 1 The trajectory of player i defined by $X_{t+1,N}^i \sim Q(x_{t,N}^i, a_{t,N}^i, m_{t,\sigma,N})$.
- 2 The trajectory defined by $X_{t+1}^i \sim Q(x_t^i, a_{t,N}^i, m_t)$.

The following estimate holds:

$$\max_{i=1,\dots,N} \mathbb{E} \left(\max_{s \leq T} \|X_s^i - X_{s,N}^i\|_\infty \right) \leq \frac{L_Q T |\mathcal{X}| \exp(T(\|Q\|_\infty + L_Q + 1))}{\sqrt{N}}$$

Some remarks

- In continuous time mean field games, the complementary approach of studying the limit behaviour of equilibria of N player games as $N \rightarrow +\infty$ has been developed by Bardi (2012) for the linear-quadratic case and by Lasry and Lions (2007) and Feleqi (2013) for games with several populations of players and ergodic payoffs.

Some remarks(cont.)

- This construction is "robust" with respect to the number of players: players can "play well" even if they do not know the exact number of players.

Some remarks(cont.)

- This construction is "robust" with respect to the number of players: players can "play well" even if they do not know the exact number of players.
- In general, the set of Nash equilibria of N player might contain equilibria that depend on N as in the driving game.

Some remarks(cont.)

- This construction is "robust" with respect to the number of players: players can "play well" even if they do not know the exact number of players.
- In general, the set of Nash equilibria of N player might contain equilibria that depend on N as in the driving game.

Driving game

- N players have to choose whether to drive on the left or on the right in a two-way road.

Driving game

- N players have to choose whether to drive on the left or on the right in a two-way road.
- Some equilibria: everyone on the left, everyone on the right, if N is even, everyone on the left; N odd, everyone on the right.

Driving game

- N players have to choose whether to drive on the left or on the right in a two-way road.
- Some equilibria: everyone on the left, everyone on the right, if N is even, everyone on the left; N odd, everyone on the right.
- The N player game has $2^{\lfloor N/2 \rfloor}$ equilibria, while the game with infinitely many players has only two.

The discounted N -player game

The λ - discounted N player game is the game with payoff:

$$J_N^{\lambda,i}(x, m_0, \pi^i, \pi^{-i}) := \mathbb{E}_{\pi}^Q \left[\sum_{t=1}^{\infty} (1 - \lambda)^{t-1} \ell(x_{t,N}^i, a_{t,N}^i, m_{t,N}) \mid x_0^i = x \right].$$

for $\lambda \in (0, 1]$.

Discounted payoff: One player

For the discounted case, one can define a mean field equilibrium as follows:
The value of the one-player game Γ_m^λ , ($m \in \mathcal{P}(\mathcal{X})$ is fixed) with payoff:

Discounted payoff: One player

For the discounted case, one can define a mean field equilibrium as follows:
 The value of the one-player game Γ_m^λ , ($m \in \mathcal{P}(\mathcal{X})$ is fixed) with payoff:

$$V_\lambda(x, m) := \sup_{\sigma} \mathbb{E}^Q \left[\sum_{t=1}^{\infty} (1 - \lambda)^{t-1} \ell(x_t, \sigma(x_t), m) \right]$$

satisfies:

Discounted payoff: One player

For the discounted case, one can define a mean field equilibrium as follows: The value of the one-player game Γ_m^λ , ($m \in \mathcal{P}(\mathcal{X})$ is fixed) with payoff:

$$V_\lambda(x, m) := \sup_{\sigma} \mathbb{E}^Q \left[\sum_{t=1}^{\infty} (1 - \lambda)^{t-1} \ell(x_t, \sigma(x_t), m) \right]$$

satisfies:

$$V_\lambda(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + (1 - \lambda) \sum_{y \in \mathcal{X}} V_\lambda(y, m) Q(x, a, m)(y) \right\}.$$

Stationary mean field equilibrium

The mean field equilibrium in this case is a fixed point of the maps $\Psi_\lambda : \mathcal{P}(\mathcal{X}) \mapsto \Sigma_s$ and $\Phi_\lambda : \Sigma_s \rightsquigarrow \mathcal{P}(\mathcal{X})$:

$$\begin{aligned}
 m &\rightsquigarrow \text{Optimal stationary strategies in } \Gamma_m^\lambda, \\
 \sigma \in \Sigma_s &\rightsquigarrow \text{Inv. dist. of the MC with transition } Q(\cdot, \sigma(\cdot), m)
 \end{aligned}$$

Existence

Proposition

$\Phi_\lambda \circ \Psi_\lambda$ has a fixed point, i.e. there exists a stationary mean field equilibrium in the following cases:

- If for every stationary strategy σ and all $m \in \mathcal{P}(\mathcal{X})$, the Markov chain with transition law $Q(\cdot, \cdot, \sigma(x), m)$ has a unique stationary distribution.

Existence

Proposition

$\Phi_\lambda \circ \Psi_\lambda$ has a fixed point, i.e. there exists a stationary mean field equilibrium in the following cases:

- If for every stationary strategy σ and all $m \in \mathcal{P}(\mathcal{X})$, the Markov chain with transition law $Q(\cdot, \cdot, \sigma(x), m)$ has a unique stationary distribution.
- The transitions are independent of the state distribution, i.e.

$$Q(x, a, m)(y) =: Q(x, a)(y), \quad \forall (x, y, a, m).$$

The short-stage equilibrium

We adapt an idea recently introduced by Neyman(2013) and Cardaliaguet et al.(2013) to our model. The aim (informally) is to construct an approximation by Friedman/Fleming discretization of a game in continuous time.

The short-stage equilibrium

We adapt an idea recently introduced by Neyman(2013) and Cardaliaguet et al.(2013) to our model. The aim (informally) is to construct an approximation by Friedman/Fleming discretization of a game in continuous time.

Consider a function $\mu : \mathcal{X} \times \mathcal{X} \times A \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ bounded and such that, for all (x, a, m) :

$$\mu(x, y, a, m) \geq 0, \quad x \neq y, \quad \mu(x, x, a, m) = - \sum_{y \neq x} \mu(x, y, a, m).$$

Short-stage, one-player game

Let $\rho > 0$. The payoff in continuous time we want to approximate is:

$$\int_0^{\infty} e^{-\rho t} \ell(x_t, \sigma(x_t), m) dt$$

where $(x_t)_{t \geq 0}$ is a continuous time Markov chain whose transition semigroup has generator $\mu(\cdot, \sigma(\cdot), m)$.

The players choose their actions at times $\{\delta, 2\delta, 3\delta, \dots\}$

The players choose their actions at times $\{\delta, 2\delta, 3\delta, \dots\}$
The value $V_{\rho, \delta}$ for the one-player game satisfies:

The players choose their actions at times $\{\delta, 2\delta, 3\delta, \dots\}$

The value $V_{\rho, \delta}$ for the one-player game satisfies:

$$\rho V_{\rho, \delta}(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + e^{-\rho\delta} \sum_{y \in \mathcal{X}} \mu(x, y, a, m) V_{\rho, \delta}(y, m) \right\}$$

The short-stage limit

Proposition

The equation

$$\rho f(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + \sum_{y \in \mathcal{X}} \mu(x, y, a, m) f(y, m) \right\}$$

has a unique fixed point, denoted V_ρ . Moreover, $V_{\rho\delta}^\delta \rightarrow V_\rho$ uniformly as $\delta \rightarrow 0$.

The short-stage limit

Proposition

The equation

$$\rho f(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + \sum_{y \in \mathcal{X}} \mu(x, y, a, m) f(y, m) \right\}$$

has a unique fixed point, denoted V_ρ . Moreover, $V_{\rho\delta}^\delta \rightarrow V_\rho$ uniformly as $\delta \rightarrow 0$.

Proof.

Use the stationary strategy σ^ρ given by this equation in the game with stage δ . □

The limit mass equation

Now let σ_δ be a fixed stationary strategy for the game with stage δ and let $m' \in \mathcal{P}(\mathcal{X})$. Let $L[\sigma_\delta, m'] \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be defined by

$$L[\sigma_\delta, m']_{x,y} = \mu(x, y, \sigma_\delta(x), m').$$

The limit mass equation

Now let σ_δ be a fixed stationary strategy for the game with stage δ and let $m' \in \mathcal{P}(\mathcal{X})$. Let $L[\sigma_\delta, m'] \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be defined by

$$L[\sigma_\delta, m']_{x,y} = \mu(x, y, \sigma_\delta(x), m').$$

The transition matrix associated to Q_δ is then

$$\mathbb{I} + \delta L[\sigma_\delta, m'].$$

The limit mass equation

Now let σ_δ be a fixed stationary strategy for the game with stage δ and let $m' \in \mathcal{P}(\mathcal{X})$. Let $L[\sigma_\delta, m'] \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be defined by

$$L[\sigma_\delta, m']_{x,y} = \mu(x, y, \sigma_\delta(x), m').$$

The transition matrix associated to Q_δ is then

$$\mathbb{I} + \delta L[\sigma_\delta, m'].$$

For σ, m' given, the associated invariant distribution m must solve

$$\delta L[\sigma, m'] \cdot m = 0 \iff L[\sigma, m'] \cdot m = 0$$

The limit system

Combining the two limit equations we obtain the following:

The limit system

Combining the two limit equations we obtain the following:

Definition

A **limit stationary mean field equilibrium** is a pair $(\sigma, m) \in \Sigma_s \times \mathcal{P}(\mathcal{X})$ such that σ is the stationary strategy derived from

$$\rho V^\rho(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + \sum_{y \in \mathcal{X}} \mu(x, y, a, m) V^\rho(y, m) \right\}$$

The limit system

Combining the two limit equations we obtain the following:

Definition

A **limit stationary mean field equilibrium** is a pair $(\sigma, m) \in \Sigma_s \times \mathcal{P}(\mathcal{X})$ such that σ is the stationary strategy derived from

$$\rho V^\rho(x, m) = \max_{a \in A} \left\{ \ell(x, a, m) + \sum_{y \in \mathcal{X}} \mu(x, y, a, m) V^\rho(y, m) \right\}$$

and m solves:

$$L[\sigma^\rho, m] \cdot m = 0.$$

Approximation

Theorem

(M.2013). For every $\epsilon > 0$ there exists $\delta_0 > 0$ and $N_0 \in \mathbb{N}$ such that, for all $\delta < \delta_0$ and $N > N_0$, the strategy provided by the limit stationary mean field equilibrium is a 2ϵ -Nash equilibrium of the discounted mean field game with discount factor $\lambda = 1 - e^{-\rho\delta}$.

Proof.

Collect everything:

Proof.

Collect everything:

- Choose δ_0 first, from the uniform convergence of the value functions.

Proof.

Collect everything:

- Choose δ_0 first, from the uniform convergence of the value functions.
- Now choose K_0 such that

$$(1 - \rho\delta_0)^{K_0} \|\ell\|_\infty < \epsilon/2.$$

Proof.

Collect everything:

- Choose δ_0 first, from the uniform convergence of the value functions.
- Now choose K_0 such that

$$(1 - \rho\delta_0)^{K_0} \|\ell\|_\infty < \epsilon/2.$$

- Finally, take N_0 as in the bound derived for the error in the game with K_0 stages.

□

Example: Online hotel booking

Example

Consider the industry of online hotel booking, where many firms offer accommodation. In this case:

- i) The state variable is the reputation of each firm.*

Example: Online hotel booking

Example

Consider the industry of online hotel booking, where many firms offer accommodation. In this case:

- i) The state variable is the reputation of each firm.*
- ii) The actions are the adjustment of the offers.*

Example: Online hotel booking

Example

Consider the industry of online hotel booking, where many firms offer accommodation. In this case:

- i) The state variable is the reputation of each firm.*
- ii) The actions are the adjustment of the offers.*
- iii) Each firm's reputation changes randomly, depending on the current reputation and the actions.*

Example: Online hotel booking

Example

Consider the industry of online hotel booking, where many firms offer accommodation. In this case:

- i) The state variable is the reputation of each firm.*
- ii) The actions are the adjustment of the offers.*
- iii) Each firm's reputation changes randomly, depending on the current reputation and the actions.*
- iv) The firms payoff depends on the average reputation of the others.*

Example: Online hotel booking

Example

Consider the industry of online hotel booking, where many firms offer accommodation. In this case:

- i) The state variable is the reputation of each firm.*
- ii) The actions are the adjustment of the offers.*
- iii) Each firm's reputation changes randomly, depending on the current reputation and the actions.*
- iv) The firms payoff depends on the average reputation of the others.*
- v) Frequent interaction is desirable in this example!*

Some remarks

- Mean field games provide an extremely simple strategy that does not need to keep track of the other players.

Some remarks

- Mean field games provide an extremely simple strategy that does not need to keep track of the other players.
- However, the mean field equilibrium might not be unique! Unless the players agree somehow on which equilibrium to play, it is hard to predict anything.

A possible way out (Repeated driving game)

Example

Consider the repeated version of the driving game with N players with the following adaptation mechanism:

- i) *Each player chooses left or right with probability $\frac{1}{2}$ on the first stage.*

A possible way out (Repeated driving game)

Example

Consider the repeated version of the driving game with N players with the following adaptation mechanism:

- i) Each player chooses left or right with probability $\frac{1}{2}$ on the first stage.*
- ii) On the second stage, observing the realizations of the first stage, each player looks at everyone's choice (and recalls its own) and imitates the choice of the majority.*

A possible way out (Repeated driving game)

Example

Consider the repeated version of the driving game with N players with the following adaptation mechanism:

- i) Each player chooses left or right with probability $\frac{1}{2}$ on the first stage.*
- ii) On the second stage, observing the realizations of the first stage, each player looks at everyone's choice (and recalls its own) and imitates the choice of the majority.*
- iii) Thus, from stage three, the players follow a mean field equilibrium.*

Possible extensions

- How do players find the mean field equilibrium? Incorporate learning mechanisms.

Possible extensions

- How do players find the mean field equilibrium? Incorporate learning mechanisms.
- Validation of this model in applications: Several economic applications in the paper of Adlakha, Johari, Weintraub(2012), applications for dynamic auctions by Iyer, Johari, Sundararajan (2011).

Thank you!